

A complete class of optimal quantum tests for unknown unitary dynamics

Lluis Eriksson

Independent researcher, Sweden

15 October 2026 — Research draft

Abstract

We consider repeated forward access to an unknown quantum channel. The null hypothesis is an arbitrary fixed unitary channel; the alternative is a memoryless channel whose input and output rotations are unknown but fixed across all uses. We prove a complete-class theorem: every physical positive tester is dominated, in orbit-averaged alternative acceptance and worst-case null acceptance, by a classical randomization of explicitly realizable parallel tests. The tests are indexed by partitions of the query number, prepare a maximally mixed representation register, entangle the multiplicity register with a reference, and measure return to that entanglement. Consequently, for any compact family of noise channels, the exact finite-query minimax problem reduces to a classical optimization over partitions. For one specified noise channel the complete error tradeoff is a minimum of explicit Choi-polynomial invariants. Adaptation and any larger positive-tester class containing parallel protocols give the same optimum. We derive the channel-independent optimal two-query test and its failure probability in terms of Choi purity and nonunitarity. For complete depolarization we obtain the exact error at every dimension, query number and parallel reference dimension, recovering the known exponential benefit of multiplicity entanglement. We also give an error exponent for orthogonal random-unitary noise. The result requires both unknown rotations and memoryless uses; it does not establish parallel optimality for calibrated noise or arbitrary channel discrimination.

Keywords: quantum hypothesis testing; unknown unitary channels; Schur–Weyl duality; quantum testers; minimax risk; reference dimension.

1 Introduction

An unknown unitary transformation can scramble the coordinates of a quantum state while preserving reversibility. A noisy channel can change both coordinates and purity. Distinguishing these possibilities is therefore a different task from checking whether a device implements one prescribed gate. If neither the input nor the output orientation of the noise is calibrated, a natural performance criterion is the worst case over those orientations.

This paper gives an exact reduction of that problem at finite query number. The reduction holds simultaneously for every memoryless noise channel, rather than only for an example chosen in advance. Once the unknown rotations have been accounted for, every admissible test can be replaced by a lottery over a finite family of parallel tests. The resulting weights are nonnegative and sum to the test’s Haar-averaged null acceptance. This is stronger than exhibiting one good probe or an asymptotic scaling: it identifies a complete class of decision procedures for every rotation-invariant alternative family.

The key point is a joint use of two symmetries. Independent collective unitary rotations on the input and output identify representation blocks. Simultaneous permutation of the identical channel uses then singles out a one-dimensional invariant vector in each matched

pair of multiplicity spaces. On the support of the unitary null, every twirled noise operator is scalar on the resulting blocks. Crucially, each scalar block corresponds to a physical parallel test. Positivity outside the null support completes the converse without requiring the detailed normalization constraints of an adaptive circuit.

Relation to prior work. Quantum testers, convex duality, symmetry reductions and minimax process discrimination are established tools; see Nakahira and Kato [1]. Chiribella and Ebler [2] introduced multiplicity-entangled probes and coherent superpositions of singlet groupings for causal identification. Their Appendix B also identifies balanced partitions through dimension/multiplicity ratios and, in Eqs. (43)–(44), the finite Schmidt-rank truncation by $\min(K, m_\lambda)$. Neither these constructions nor the doubled depolarizing exponent are claimed here as new.

Parallel optimality itself also has substantial precedents. Hayashi [4] proves it for compact-group unitary estimation. Chen et al. [5] and Hayashi et al. [6] study unitary symmetry testing, including positive-tester converses and an exact max-relative-entropy formula. Their hypotheses concern distributions of unitary operations and subgroup membership. Here the alternative consists of tensor powers of general, potentially nonunital channels, independently rotated at their input and output. The explicit simultaneous reduction for all such alternatives is the proposed contribution. The proof uses the additional permutation symmetry of the memoryless channel Choi operator.

Unitarity estimation [7], testing against a known target under depolarizing noise [8], and comparison of unknown unitaries [9] address related but different statistical promises. The null in non-mixed-unitarity testing [10] includes every mixed-unitary channel; our null includes only a single unknown unitary repeated throughout the experiment. These null families must not be interchanged. Resource constraints in subgroup testing have recently been studied in [11]; the reference-dimension result below concerns complete depolarization and a different null.

Earlier work by the present author used algebraic query support for restricted unitary families [12], and established a minimax parallel formula for Schur noise with an unknown phase carrier together with an adaptive advantage after calibration [13]. The present result handles arbitrary CPTP channels under full, independent input and output rotations. It does not remove the calibration distinction established in that work. The literature comparison is a priority assessment for this draft, not an independent confirmation of novelty.

2 Statistical model and notation

Let $d \geq 2$, $n \geq 1$, and let $\mathcal{E} : \mathcal{L}(\mathbb{C}^d) \rightarrow \mathcal{L}(\mathbb{C}^d)$ be completely positive and trace preserving (CPTP). Write $\mathcal{U}_W(X) = W X W^\dagger$. The two hypotheses are

$$H_0 : \quad \mathcal{U}_W \quad \text{for an arbitrary } W \in U(d), \quad (1)$$

$$H_1 : \quad \mathcal{U}_L \circ \mathcal{E} \circ \mathcal{U}_R \quad \text{for } \mathcal{E} \in \mathcal{C}, \quad L, R \in U(d), \quad (2)$$

where $\mathcal{C}$ is a nonempty compact family of CPTP maps. The same W , or the same triple $(\mathcal{E}, L, R)$, is used at all n calls. The calls under H_1 are memoryless, so their joint Choi operator is a tensor power. There is no access to inverse gates, controlled versions of an unspecified dilation, or the environment.

We use the unnormalized Choi convention

$$J_{\mathcal{E}} = \sum_{i,j=1}^d \mathcal{E}(|i\rangle\langle j|) \otimes |i\rangle\langle j|, \quad \text{Tr}_O J_{\mathcal{E}} = I_I, \quad \text{Tr } J_{\mathcal{E}} = d. \quad (3)$$

All output legs are placed before all input legs. In this ordering, $J_{\mathcal{E}}^{\otimes n}$ always means the tensor product followed by that fixed regrouping. For a unitary, $J_{\mathcal{U}_W} = |W\rangle\rangle\langle\langle W|$, with vectorization in the displayed basis.

The outcome called *accept* accepts H_0 . Every ordinary adaptive protocol with n calls has a positive tester effect T such that its acceptance probability on a channel $\mathcal{F}$ is

$$a_T(\mathcal{F}) = \text{Tr}(T J_{\mathcal{F}}^{\otimes n}). \quad (4)$$

The converse will use only $T \succeq 0$ and $0 \leq a_T(\mathcal{U}_W) \leq 1$. It therefore also applies to any more general physically normalized positive-tester model, including indefinite-order models with this multilinear Choi representation. This is not a claim that every positive matrix is a physical protocol.

Define the errors

$$\alpha(T) = 1 - \inf_{W \in U(d)} a_T(\mathcal{U}_W), \quad (5)$$

$$\beta_{\mathcal{C}}(T) = \sup_{\mathcal{E} \in \mathcal{C}} \sup_{L, R \in U(d)} a_T(\mathcal{U}_L \circ \mathcal{E} \circ \mathcal{U}_R). \quad (6)$$

For a strategy class $\mathcal{S}$, let $\beta_{n, \mathcal{C}}^{\mathcal{S}}(\alpha)$ be the infimum of the latter error subject to $\alpha(T) \leq \alpha$. Randomized decisions are allowed. Haar measures below have mass one.

2.1 Schur blocks and explicit invariants

Let $\mathcal{Y}_{n,d}$ be the partitions $\lambda \vdash n$ with at most d rows. Schur–Weyl duality gives

$$\mathcal{H}_n = (\mathbb{C}^d)^{\otimes n} = \bigoplus_{\lambda \in \mathcal{Y}_{n,d}} V_{\lambda} \otimes M_{\lambda}, \quad q_{\lambda} = \dim V_{\lambda}, \quad m_{\lambda} = \dim M_{\lambda}. \quad (7)$$

Write Π_{λ} for the isotypic projector. If P_{π} permutes the n physical factors, define

$$S_n = \frac{1}{n!} \sum_{\pi \in S_n^{\text{perm}}} P_{\pi}^O \otimes P_{\pi}^I, \quad D_{\lambda} = (\Pi_{\lambda}^O \otimes \Pi_{\lambda}^I) S_n. \quad (8)$$

Here S_n^{perm} denotes the permutation group, to distinguish it from the averaging projector S_n . The permutation matrices in the computational basis are real. Input Schur spaces are identified with the conjugate spaces in the vectorization convention. In compatible Schur bases,

$$D_{\lambda} = I_{V_{\lambda}^O} \otimes I_{V_{\lambda}^I} \otimes \Phi_{m_{\lambda}}, \quad \Phi_m = |\phi_m\rangle\langle\phi_m|, \quad |\phi_m\rangle = \frac{1}{\sqrt{m}} \sum_{a=1}^m |a, a\rangle. \quad (9)$$

Thus the D_{λ} are orthogonal projectors of rank q_{λ}^2 .

The quantities that determine the answer are

$$b_{\lambda}(\mathcal{E}) = \frac{\text{Tr}(D_{\lambda} J_{\mathcal{E}}^{\otimes n})}{q_{\lambda} m_{\lambda}}. \quad (10)$$

They are explicitly defined degree- n polynomials in the Choi entries. They lie in $[0, 1]$ and are invariant under independent unitary rotations at the input and output, as also follows from the physical realization below. Their definition does not entail a polynomial-time evaluation algorithm in n and d .

3 A complete class and the exact minimax formula

Theorem 3.1 (Simultaneous parallel domination). *For every physical positive tester T , put*

$$c = \int a_T(\mathcal{U}_W) dW, \quad w_\lambda = \frac{m_\lambda}{q_\lambda} \text{Tr}(TD_\lambda). \quad (11)$$

Then $w_\lambda \geq 0$ and $\sum_\lambda w_\lambda = c \leq 1$. There is a parallel randomized test $\tilde{T}$, depending only on these weights, such that

$$a_{\tilde{T}}(\mathcal{U}_W) = c \quad \text{for every } W, \quad (12)$$

$$a_{\tilde{T}}(\mathcal{U}_L \circ \mathcal{E} \circ \mathcal{U}_R) = \sum_\lambda w_\lambda b_\lambda(\mathcal{E}) \leq \int a_T(\mathcal{U}_{L'} \circ \mathcal{E} \circ \mathcal{U}_{R'}) dL' dR' \quad (13)$$

for every CPTP channel $\mathcal{E}$ and every L, R .

The inequality is with the original test's *orbit average*, not its performance at each particular orientation. Since an average is at most a supremum, the theorem gives simultaneous domination for the worst-case risk used here. At the null, the constant c is at least the original test's minimum acceptance.

Corollary 3.2 (Robust alternatives). *For $0 \leq \alpha \leq 1$, let $\Delta(\mathcal{Y}_{n,d})$ denote the probability simplex on $\mathcal{Y}_{n,d}$. Then*

$$\boxed{\beta_{n,\mathcal{C}}^S(\alpha) = (1 - \alpha) \min_{p \in \Delta(\mathcal{Y}_{n,d})} \sup_{\mathcal{E} \in \mathcal{C}} \sum_{\lambda \in \mathcal{Y}_{n,d}} p_\lambda b_\lambda(\mathcal{E}).} \quad (14)$$

This holds for parallel strategies, all causal strategies, and any physically normalized positive-tester class $\mathcal{S}$ between them and the positivity relaxation used in Theorem 3.1. A minimizing p is implemented by selecting a partition with probability p_λ , performing its parallel test, and accepting a successful outcome with probability $1 - \alpha$.

Proof. For $\alpha < 1$, an admissible T has $c \geq 1 - \alpha > 0$. Set $p_\lambda = w_\lambda/c$ in Theorem 3.1. Nonnegativity of every b_λ gives the lower bound in (14). The stated construction attains the upper bound. The objective is continuous on a compact simplex, so a minimizer exists. At $\alpha = 1$, always rejecting gives zero. $\square$

For finitely many canonical noise models $\mathcal{E}_1, \dots, \mathcal{E}_k$, the remaining optimization is the ordinary linear program

$$\min_{p,t} t \quad \text{subject to} \quad p \geq 0, \quad \sum_\lambda p_\lambda = 1, \quad \sum_\lambda p_\lambda b_\lambda(\mathcal{E}_j) \leq t \quad (1 \leq j \leq k). \quad (15)$$

The same theorem applies to any compact alternative set invariant under independent input and output rotations, including distance-based promises invariant under those rotations. Computing its supremum over channels can remain difficult. If the alternative contains a unitary channel, the optimum is necessarily $1 - \alpha$; a nontrivial uniform testing problem requires a separated alternative.

Corollary 3.3 (A specified channel up to rotations). *For $\mathcal{C} = \{\mathcal{E}\}$,*

$$\boxed{\beta_{n,\mathcal{E}}(\alpha) = (1 - \alpha) \min_{\lambda \in \mathcal{Y}_{n,d}} b_\lambda(\mathcal{E}).} \quad (16)$$

One partition, with a reference of dimension m_λ , suffices. The entire optimal error tradeoff is linear in $1 - \alpha$.

4 Proof of the complete-class theorem

4.1 The support of the unitary null

Define

$$M_0 = \int |W^{\otimes n}\rangle\langle W^{\otimes n}| dW. \quad (17)$$

In (7), $W^{\otimes n} = \bigoplus_{\lambda} W_{\lambda} \otimes I_{M_{\lambda}}$. Schur orthogonality eliminates cross terms between inequivalent $U(d)$ representations and sends $|W_{\lambda}\rangle\langle W_{\lambda}|$ to I/q_{λ} . The vectorization of $I_{M_{\lambda}}$ has squared norm m_{λ} . Hence

$$M_0 = \sum_{\lambda \in \mathcal{Y}_{n,d}} \frac{m_{\lambda}}{q_{\lambda}} D_{\lambda}. \quad (18)$$

The identity (9) follows independently by averaging the diagonal permutation action: the trivial representation occurs in $M_{\lambda} \otimes M_{\mu}$ precisely when $\lambda = \mu$, and then with multiplicity one. This uses the real self-duality of symmetric-group irreducibles. In particular $\text{supp } M_0 = \bigoplus_{\lambda} \text{ran } D_{\lambda}$.

4.2 The noise block lemma

Let

$$B_{\mathcal{E}} = \int J_{\mathcal{U}_L \circ \mathcal{E} \circ \mathcal{U}_R}^{\otimes n} dL dR. \quad (19)$$

In the chosen convention $J_{\mathcal{U}_L \circ \mathcal{E} \circ \mathcal{U}_R} = (L \otimes R^T) J_{\mathcal{E}} (L \otimes R^T)^{\dagger}$. Both Haar twirls are therefore collective unitary twirls on the two sides.

Lemma 4.1. *For every CPTP $\mathcal{E}$,*

$$B_{\mathcal{E}} = \sum_{\lambda \in \mathcal{Y}_{n,d}} \frac{m_{\lambda}}{q_{\lambda}} b_{\lambda}(\mathcal{E}) D_{\lambda} + B_{\mathcal{E}}^{\perp}, \quad B_{\mathcal{E}}^{\perp} \succeq 0, \quad B_{\mathcal{E}}^{\perp} D_{\lambda} = 0. \quad (20)$$

Proof. The independent twirls make $B_{\mathcal{E}}$ block diagonal in the labels (λ, μ) and proportional to the identity on $V_{\lambda}^O \otimes \overline{V}_{\mu}^I$. Moreover, $J_{\mathcal{E}}^{\otimes n}$ is invariant under conjugation by $P_{\pi}^O \otimes P_{\pi}^I$, because these operations jointly permute identical Choi factors. The twirls preserve this invariance.

In the (λ, λ) block, the multiplicity operator thus commutes with the diagonal symmetric-group action. Its trivial subrepresentation is the single vector $|\phi_{m_{\lambda}}\rangle$. It cannot couple that vector to a nontrivial subrepresentation. Consequently D_{λ} is a reducing subspace on which $B_{\mathcal{E}}$ is scalar. Since D_{λ} commutes with both twirls,

$$B_{\mathcal{E}} D_{\lambda} = \frac{\text{Tr}(D_{\lambda} B_{\mathcal{E}})}{q_{\lambda}^2} D_{\lambda} = \frac{\text{Tr}(D_{\lambda} J_{\mathcal{E}}^{\otimes n})}{q_{\lambda}^2} D_{\lambda} = \frac{m_{\lambda}}{q_{\lambda}} b_{\lambda}(\mathcal{E}) D_{\lambda}. \quad (21)$$

The restriction to the orthogonal complement is positive, which proves (20). $\square$

Taking the trace against $T \succeq 0$ proves

$$\text{Tr}(T B_{\mathcal{E}}) \geq \sum_{\lambda} w_{\lambda} b_{\lambda}(\mathcal{E}). \quad (22)$$

Equation (18) gives $\sum w_{\lambda} = c$. These are the only facts needed for the converse.

4.3 Realization of every block

For a selected partition λ , prepare on the n input systems and a reference $A \simeq \mathbb{C}^{m_\lambda}$ the state

$$\rho_\lambda = \frac{I_{V_\lambda}}{q_\lambda} \otimes \Phi_{M_\lambda A}, \quad (23)$$

embedded in the λ sector of $\mathcal{H}_n$. Apply the n channel calls in parallel and accept with the effect

$$F_\lambda = I_{V_\lambda^O} \otimes \Phi_{M_\lambda^O A}, \quad (24)$$

extended by zero on the other output sectors. No purification of the maximally mixed representation register needs to be retained.

The null channel acts only on V_λ , so acceptance is exactly one for every W . The input state is invariant under collective input rotations and the measurement under collective output rotations. Acceptance under $\mathcal{E}$ is therefore orientation independent. A direct Choi contraction of (23) and (24) gives the tester

$$T_\lambda = \frac{D_\lambda}{q_\lambda m_\lambda}, \quad a_{T_\lambda}(\mathcal{E}) = b_\lambda(\mathcal{E}). \quad (25)$$

For example $\text{Tr}(T_\lambda M_0) = 1$, checking all normalization factors. Choosing this test with probability w_λ and otherwise rejecting yields Theorem 3.1.

Remark 4.2 (Why both rotations matter). The proof does not symmetrize an adaptive circuit and assert that it becomes parallel. It first proves a positive-operator decomposition, then realizes each surviving block in parallel. An output twirl alone need not make the multiplicity/input operator scalar on D_λ . Likewise, correlated noise need not have the required permutation symmetry or tensor-power Choi operator. Neither case is covered by the theorem.

5 The universal optimal two-query test

At $n = 1$ there is only one sector and $b_{(1)}(\mathcal{E}) = \text{Tr } J_\mathcal{E}/d = 1$: one call has no power for this minimax problem. At $n = 2$, write $+$ and $-$ for the symmetric and antisymmetric sectors. Their multiplicities are one and their representation dimensions are $q_\pm = d(d \pm 1)/2$.

Theorem 5.1. *For every d -dimensional CPTP channel, set*

$$s = \text{Tr}[\mathcal{E}(I)^2], \quad p = \text{Tr}(J_\mathcal{E}^2). \quad (26)$$

Then

$$b_\pm(\mathcal{E}) = \frac{d^2 \pm (s + d) + p}{2d(d \pm 1)}, \quad (27)$$

$$\beta_{2,\mathcal{E}}(\alpha) = (1 - \alpha) \frac{d^2 - d - s + p}{2d(d - 1)}. \quad (28)$$

Antisymmetric probes have antecedents in unambiguous unitary comparison [3]. In the present two-sided rotation model, the antisymmetric test is optimal simultaneously for every $\mathcal{E}$. It uses the normalized antisymmetric projector as input and checks antisymmetry at the output, with no retained reference system.

Proof. Let F_O, F_I swap the two output and input registers. Since $m_\pm = 1$,

$$D_\pm = \frac{1}{4}(I \pm F_O)(I \pm F_I). \quad (29)$$

The four contractions against $J_{\mathcal{E}}^{\otimes 2}$ are respectively d^2 , s , d , and p , proving (27). Subtracting gives

$$b_+(\mathcal{E}) - b_-(\mathcal{E}) = \frac{ds - p}{d(d^2 - 1)} \geq 0. \quad (30)$$

Indeed $\mathcal{E}(I) \succeq 0$ has trace d , so $s \geq d$, while $J_{\mathcal{E}} \succeq 0$ has trace d , so $p \leq d^2$. Apply Corollary 3.3. $\square$

For qubits, write the Bloch action as $x \mapsto Ax + t$, with a real 3×3 matrix A . Then

$$\beta_{2,\mathcal{E}}(\alpha) = (1 - \alpha) \frac{1 + \|A\|_F^2 - \|t\|_2^2}{4}. \quad (31)$$

This is an operational meaning of a familiar combination of purity and nonunitarity invariants; the SWAP contraction itself is standard. For amplitude damping with damping probability γ , the value at $\alpha = 0$ is $1 - \gamma$. For a replacement channel $X \mapsto \text{Tr}(X)\sigma$, it is $(1 - \text{Tr} \sigma^2)/2$, so a pure-state replacement is perfectly distinguishable using two queries. For complete depolarization, it is $(d - 1)/(2d)$ in dimension d .

The formula also immediately solves any two-query robust alternative: take the supremum of (28) over its canonical channels. No optimization over a different input state for each channel is necessary.

6 Exact reference cost for complete depolarization

Let $\mathcal{D}(X) = \text{Tr}(X)I/d$. Then $J_{\mathcal{D}} = I/d$, and

$$b_{\lambda}(\mathcal{D}) = \frac{q_{\lambda}}{d^n m_{\lambda}}. \quad (32)$$

We refine this benchmark by bounding the dimension $K \geq 1$ of the quantum reference retained outside the n parallel calls. Probe entanglement among the n inputs is unrestricted. This constraint is not the incoherent-access model of [7], nor a bound on all working memory in an adaptive circuit.

Theorem 6.1 (Finite reference dimension). *Write $n = ad + r$, where $0 \leq r < d$, and define*

$$\lambda_* = (\underbrace{a+1, \dots, a+1}_r, \underbrace{a, \dots, a}_{d-r}), \quad q_* = \binom{d}{r}, \quad m_* = \dim M_{\lambda_*}. \quad (33)$$

For parallel tests with reference dimension at most K ,

$$\beta_{n,\mathcal{D}}^{\text{par},K}(\alpha) = (1 - \alpha) \min_{\lambda \in \mathcal{Y}_{n,d}} \frac{q_{\lambda}}{d^n \min(K, m_{\lambda})} = (1 - \alpha) \frac{\binom{d}{r}}{d^n \min(K, m_*)}. \quad (34)$$

With unrestricted reference, this is also the optimum for all the tester classes in Corollary 3.2.

Proof. First twirl an arbitrary parallel input by a collective unitary. This does not change the alternative output. Its null acceptance, for a fixed measurement and any W , becomes the original Haar-averaged null acceptance, which is at least $1 - \alpha$. The input can thus be written

$$\rho = \bigoplus_{\lambda} \frac{I_{V_{\lambda}}}{q_{\lambda}} \otimes X_{\lambda}, \quad X_{\lambda} \succeq 0, \quad \sum_{\lambda} \text{Tr} X_{\lambda} = 1. \quad (35)$$

The reference marginal is $\rho_A = \sum_{\lambda} \text{Tr}_{M_{\lambda}} X_{\lambda}$. Twirling the measurement at the output changes neither acceptance probability, and gives blocks $I_{V_{\lambda}} \otimes Q_{\lambda}$, with $0 \preceq Q_{\lambda} \preceq I$.

Table 1: Exact depolarizing failure probabilities for qubits at zero false positives. The reference dimension is K ; the last column is the unrestricted optimum, also valid against adaptive testers.

n	q_*	m_*	$K = 1$	$K \geq m_*$
2	1	1	1/4	1/4
3	2	2	1/4	1/8
4	1	2	1/16	1/32
5	2	5	1/16	1/80
6	1	5	1/64	1/320
8	1	14	1/256	1/3584

For every positive X on $\mathbb{C}^m \otimes \mathbb{C}^K$, the Schmidt-rank inequality gives

$$X \preceq \min(m, K) I_m \otimes \text{Tr}_{\mathbb{C}^m} X. \quad (36)$$

For a rank-one X this is the Cauchy–Schwarz inequality on its Schmidt expansion; summing a rank-one decomposition proves the general case. In particular,

$$X_\lambda \preceq \min(m_\lambda, K) I_{M_\lambda} \otimes \rho_A. \quad (37)$$

The alternative acceptance consequently satisfies

$$\frac{1}{d^n} \sum_\lambda q_\lambda \text{Tr}[Q_\lambda(I_{M_\lambda} \otimes \rho_A)] \geq \min_\lambda \frac{q_\lambda}{d^n \min(K, m_\lambda)} \sum_\lambda \text{Tr}(Q_\lambda X_\lambda) \quad (38)$$

$$\geq (1 - \alpha) \min_\lambda \frac{q_\lambda}{d^n \min(K, m_\lambda)}. \quad (39)$$

Equality is attained in any selected sector by using I_{V_λ}/q_λ , a maximally entangled state of Schmidt rank $\min(K, m_\lambda)$ in $M_\lambda \otimes A$, its return projector, and acceptance thinning by $1 - \alpha$.

It remains to optimize the partition. The standard dimension formulas imply

$$\frac{q_\lambda}{m_\lambda} = \frac{\prod_{i=1}^d (\lambda_i + d - i)!}{n! \prod_{j=0}^{d-1} j!}. \quad (40)$$

Moving a box from a longer row to a shorter row, while keeping a partition and with row lengths differing by at least two, strictly decreases the numerator. Repeated balancing reaches λ_* . This is the balanced-ratio property also proved in [2, Appendix B].

We also need $q_\lambda \geq q_*$. Every partition of n with at most d parts dominates λ_* . By the classical dominance criterion for Kostka numbers [14, Chapter I], the $U(d)$ representation V_λ contains the weight λ_* . Its distinct coordinate permutations are also weights, each with positive multiplicity. There are $\binom{d}{r}$ such permutations. On the other hand V_{λ_*} is a determinant twist of $\bigwedge^r \mathbb{C}^d$, so $q_* = \binom{d}{r}$.

Thus $q_\lambda \geq q_*$ and $q_\lambda/m_\lambda \geq q_*/m_*$. If $K \leq m_*$ the first inequality gives $q_\lambda/\min(K, m_\lambda) \geq q_*/K$. If $K \geq m_*$ the second gives $q_\lambda/\min(K, m_\lambda) \geq q_*/m_*$. Both cases prove (34). $\square$

For fixed d , Stirling’s formula applied to (40) gives

$$\frac{m_*}{q_*} = \Theta_d(d^n n^{-(d^2-1)/2}), \quad \beta_{n, \mathcal{D}}(0) = \Theta_d(n^{(d^2-1)/2} d^{-2n}). \quad (41)$$

If $K_n = \exp(\kappa n + o(n))$ with $\kappa \geq 0$, then

$$\lim_{n \rightarrow \infty} -\frac{1}{n} \log \beta_{n, \mathcal{D}}^{\text{par}, K_n}(0) = \log d + \min(\kappa, \log d). \quad (42)$$

The endpoints recover the known exponential advantage of multiplicity entanglement [2]. The finite- K identity applies the Schmidt-rank truncation of [2, Appendix B, Eqs. (43)–(44)] to the present asymmetric testing problem and establishes its corresponding converse, including query numbers not divisible by d .

7 An exponent and a stability bound

Proposition 7.1 (Orthogonal random-unitary noise). *Suppose $\mathcal{E} = \sum_{a=1}^s p_a \mathcal{U}_{U_a}$, where $p_a > 0$, $\sum p_a = 1$, and $\text{Tr}(U_a^\dagger U_b) = d\delta_{ab}$. Set $p_{\max} = \max_a p_a$. For any fixed $0 \leq \alpha < 1$,*

$$\lim_{n \rightarrow \infty} -\frac{1}{n} \log \beta_{n,\mathcal{E}}(\alpha) = -\log p_{\max}. \quad (43)$$

Proof. The channel tensor power contains the component $\mathcal{U}_{U_a}^{\otimes n}$ with probability p_a^n . This component is an admissible null, also after input and output rotations. Positivity gives the lower bound $(1 - \alpha)p_{\max}^n$ for every admissible tester.

For the upper bound, send one half of a maximally entangled d -dimensional pair through each channel call. The null output is a tensor power of the same pure state in $\mathbb{C}^{d^2}$. Accepting its full symmetric subspace has null acceptance one. For the alternative, the normalized Choi state has eigenvalues $p_1, \dots, p_s$, unaffected by the two rotations. The symmetric-projector acceptance is the complete homogeneous polynomial

$$h_n(p) = \sum_{k_1 + \dots + k_s = n} \prod_a p_a^{k_a} \leq \binom{n + s - 1}{s - 1} p_{\max}^n. \quad (44)$$

Thin acceptance by $1 - \alpha$ and take logarithms. The polynomial prefactor does not affect the exponent. $\square$

This benchmark uses the standard pure-state symmetric-subspace test. It is not asserted to be optimal at finite n . In particular it can be far from the balanced-sector optimum for complete depolarization.

Proposition 7.2 (Model perturbation). *For two canonical CPTP maps $\mathcal{E}, \mathcal{F}$,*

$$|\beta_{n,\mathcal{E}}(\alpha) - \beta_{n,\mathcal{F}}(\alpha)| \leq (1 - \alpha) \min \left\{ 1, \frac{n}{2} \|\mathcal{E} - \mathcal{F}\|_\diamond \right\}. \quad (45)$$

Proof. The tensor-power telescoping identity gives $\|\mathcal{E}^{\otimes n} - \mathcal{F}^{\otimes n}\|_\diamond \leq n \|\mathcal{E} - \mathcal{F}\|_\diamond$. For a fixed probe and effect the probability difference is at most half the trace-norm difference of the output states, and at most one. Apply this to every test (23)–(24), take minima over λ , and use (16). $\square$

8 Verification, scope and remaining questions

The results above are analytic consequences of finite-dimensional representation theory and tester positivity. Computation was used to challenge normalization, feasibility and small-parameter consequences, not to replace the all-parameter proofs. The accompanying scripts and machine-readable reports record:

- 91 channel/dimension/query cases, including random CPTP maps, unitary channels, complete depolarization, pure replacement, amplitude damping and dephasing; projectors were built from permutation characters for $(d, n) = (2, 1), (2, 2), (2, 3), (2, 4), (3, 2)$;

- exact-design twirl checks for qubits up to three queries, with maximum block residual 3.1×10^{-15} and worst negative eigenvalue magnitude 1.5×10^{-15} in the predicted dual inequality;
- 12 unrestricted two-query causal semidefinite programs, using four channels and three false-positive allowances, agreeing with (16) to within 4.5×10^{-8} ; some solver outcomes were marked numerically inaccurate and are retained as such in the report;
- 75 direct probe/channel/measurement calculations with explicit three- and four-qubit multiplicity spaces, random input/output rotations and reference dimensions one or two;
- 11,626 exact rational partition comparisons and 80,909 exact finite-reference comparisons for $2 \leq d \leq 7$, $1 \leq n \leq 24$.

The checks are reproducible but are not independent peer review or proof-assistant formalization. The numerical twirl checks use the 24-element single-qubit Clifford group only in its valid unitary-three-design range; they do not substitute a Clifford twirl for a fourth Haar moment.

The principal limitation is the statistical promise. Full uncertainty in both collective orientations makes the complete-class reduction possible. The theorem neither says that adaptation is useless for general noisy-channel discrimination nor supplies an efficient laboratory implementation of every large Schur transform. Preparation and measurements are ideal, the system dimension is known, and the noise has no temporal memory. Reference dimension in Theorem 6.1 constrains only a parallel bypass system.

For a general channel, the minimizing partition in (16) remains a finite optimization problem. We do not claim that the balanced partition is universally optimal outside the depolarizing benchmark. Determining useful closed forms for broader noise families, the computational complexity of the robust optimization, and the extent to which partial calibration restores an adaptive advantage are natural next questions. A specialist assessment of the complete-class theorem and its relationship to existing covariant testing theory remains necessary before assigning a strong novelty or impact claim.

A Evaluation without Haar integration

The central projector is

$$\Pi_\lambda = \frac{m_\lambda}{n!} \sum_{\sigma \in S_n^{\text{perm}}} \chi_\lambda(\sigma) P_\sigma. \quad (46)$$

Because diagonal permutation invariants pair only identical labels, $(\Pi_\lambda \otimes I)S_n = D_\lambda$. Consequently

$$b_\lambda(\mathcal{E}) = \frac{1}{q_\lambda(n!)^2} \sum_{\sigma, \pi \in S_n^{\text{perm}}} \chi_\lambda(\sigma) \text{Tr}[(P_\sigma^O P_\pi^O \otimes P_\pi^I) J_{\mathcal{E}}^{\otimes n}]. \quad (47)$$

This is a finite expression in permutation contractions, independent of a choice of Kraus representation. It gives a direct evaluation method for small n .

For another expression, let $\{K_a\}$ be any Kraus family and $K_a = K_{a_1} \otimes \cdots \otimes K_{a_n}$. Regard the compression $K_{a,\lambda} = \Pi_\lambda K_a \Pi_\lambda$ as an operator on $V_\lambda \otimes M_\lambda$. Projection of its vectorization onto $|\phi_{m_\lambda}\rangle$ contracts the multiplicity indices and supplies a factor $m_\lambda^{-1/2}$. Thus

$$b_\lambda(\mathcal{E}) = \frac{1}{q_\lambda m_\lambda^2} \sum_a \|\text{Tr}_{M_\lambda} K_{a,\lambda}\|_{\text{HS}}^2. \quad (48)$$

This expression makes positivity transparent, whereas the physical test proves the upper bound one.

B A single operator certificate

Let $b = \min_{\lambda} b_{\lambda}(\mathcal{E})$. Equations (18) and (20) give

$$B_{\mathcal{E}} - bM_0 \succeq 0. \quad (49)$$

For every $T \succeq 0$ satisfying $\text{Tr}(TM_0) \geq 1 - \alpha$, this certifies $\text{Tr}(TB_{\mathcal{E}}) \geq b(1 - \alpha)$. Equality is attained by the explicitly physical test in any minimizing sector. In max-relative-entropy notation, with the usual support convention,

$$b = \exp[-D_{\max}(M_0 \| B_{\mathcal{E}})], \quad D_{\max}(A \| B) = \log \inf\{t : A \preceq tB\}. \quad (50)$$

The use of this divergence for exact unitary hypothesis testing has antecedents in [5, 6]. Here the block lemma computes it for rotated memoryless CPTP alternatives and proves simultaneous attainability, including the case $b = 0$.

Reproducibility and assistance. The companion directory contains the verification scripts, their output reports, the search-coverage audit and a claim-by-claim proof audit. OpenAI Codex assisted with retrieval, mathematical development, computation and typesetting. A separate GPT-6 Astra xhigh assessment and an exact-file addendum accompany the revised manuscript. They are AI assessments from the same provider, not human peer review or an external publication decision. The present revision adds bibliographic attributions without changing the mathematical statements or proofs.

References

- [1] K. Nakahira and K. Kato, Generalized quantum process discrimination problems, *Phys. Rev. A* **103**, 062606 (2021); erratum **104**, 069902 (2021). [arXiv:2104.09759](#).
- [2] G. Chiribella and D. Ebler, Quantum speedup in the identification of cause–effect relations, *Nature Communications* **10**, 1472 (2019). [doi:10.1038/s41467-019-09383-8](#); [arXiv:1806.06459](#).
- [3] M. Sedlák and M. Ziman, Unambiguous comparison of unitary channels, *Phys. Rev. A* **79**, 012303 (2009). [arXiv:0809.4401](#).
- [4] M. Hayashi, Indefinite causal order strategy does not improve the estimation of group action, *Quantum* **9**, 1891 (2025). [arXiv:2501.09312](#).
- [5] Y.-A. Chen, C. Zhu, K. He, Y. Liu and X. Wang, Hypothesis testing of symmetry in quantum dynamics (2024). [arXiv:2411.14292](#).
- [6] M. Hayashi, Y.-A. Chen, C. Zhu and X. Wang, Predicting symmetries of quantum dynamics with optimal samples (2025). [arXiv:2502.01464](#).
- [7] K. Chen, Q. Wang, P. Long and M. Ying, Unitarity estimation for quantum channels, *IEEE Trans. Inf. Theory* **69**, 5116–5134 (2023). [arXiv:2212.09319](#).
- [8] H. Ghosh, A. Jackson and A. Datta, Unitary channel testing under a depolarizing noise assumption (2026). [arXiv:2606.10730](#).
- [9] Y. Hashimoto, A. Soeda and M. Murao, Comparison of unknown unitary channels with multiple queries (2022; revised 2026). [arXiv:2208.12519](#).
- [10] P. Ghosal, P. Halder, A. Patra and A. Sen(De), Quantum hypothesis testing of non-mixed-unitarity: A multifaceted hierarchy of quantum channel discrimination (2026). [arXiv:2609.40355](#).
- [11] Y. Cao, R. Chen, C. Zhu and X. Wang, Resource compatibility in optimal quantum symmetry testing (2026). [arXiv:2609.38134](#).

- [12] L. Eriksson, Algebraic query support for unitary oracles: Exact Hilbert laws, harmonic spectra, and general-tester bounds (2026). [AIRR ARR-2026-6WX2JF38WE87GB2M](#), v1.
- [13] L. Eriksson, Phase calibration unlocks adaptive tests for renewed quantum noise (2026). [AIRR ARR-2026-33GYTDAT5C8THB3M](#), v1.
- [14] I. G. Macdonald, *Symmetric Functions and Hall Polynomials*, second edition, Oxford University Press (1995). [doi:10.1093/oso/9780198534891.001.0001](#).